

WIDE AREA NETWORK USING INTERNET WITH HIGH QUALITY OF SERVICE

By
Richard Haney

Field of use and Background of the Invention

The invention is useful in providing wide area networking services to clients with many locations among which data, especially high volumes of data, must be sent.

10 The prior art of WANs include frame relay and point-to-point networking offered by telephone companies. One type of Wide Area Network (WAN) service provided by telephone companies is leased lines. These may be analog or digital and are provided typically by a Local Exchange Carrier (LEC) on an intraLATA basis (Local Access and Transport Area). InterLATA leased lines are also available but must be provided by an Interexchange Carrier (IXC) with the LEC providing the local loop connection.

15 Another such WAN service is known as a Virtual Private Network. A VPN is intended for use by very large organizations with multiple locations. A VPN appears to the user as if it was private leased line trunk network, but it is not. VPN services are generally arranged with an Interexchange Carrier (IXC) with the points of the network termination (locations from which data will be sent and received being identified along with the level of bandwidth required at each termination. Dedicated circuits (telephone lines) are established between each network termination and the closest capable IXC POP (Point of Presence). Connections between POPs are not dedicated but are established by routers using routing tables to route the traffic over specified high-capacity transmission facilities on a priority basis to ensure the level of service provided is adequate and equivalent to a true private network using 20 leased lines.

25 Other forms of Public Data Networks include: DDS, Switched 56 Kbps; Digital T-Carrier Systems; Digital 800 Services; X.25 Packet Switched Services; Broadband Data Networking such as Frame Relay and Cell Switching, ADSL, HDSL, Sonet, Switched Megabit Data Services, ISDN and Advanced Intelligent Networks.

30 Dataphone Digital Service (DDS) which was introduced by AT&T in 1974 and is generally end-to-end, fully digital, dedicated service provided by most carriers. DDS may be either point-to-point or multipoint. A head end Front End Processor controls all access

PATENT

to the network by polling remote devices. All communication must pass through the head end. DDS signals are carried within logical channels on T1 lines.

Switched 56 Kbps is a circuit switched (rather than dedicated line) digital service that serves the same applications as DDS although it is more cost effective for lower data 5 volumes. All the components are the same as DDS but digital local loops and digital carrier exchanges are used. The main difference over DDS is that traffic is routed using a logical address which is the equivalent of a voice telephone number. The circuit is set up, maintained and torn down much like a voice call is switched and pricing is similar. The cost is sensitive to distance, duration, time of day and day of the year.

10 Digital T-carrier systems (including fractional T1 service) are dedicated links carry digital data over multiple logical channels on a single physical communication circuit with the logical channels established by time division multiplexing.

Digital 800 service was introduced in 1994 by AT&T and is intended for medium to high volume customers subscribing to high volume 800 service offerings.

15 X.25 packet switching was invented in the early 60's and was implemented on ARPANET in 1971. X.25 is a dial up service as is ISDN and Switched 56/64 Kbps WANS, and, as such, is not suitable for dedicated WANs such as the WANs in the AlterWAN™ network genus of the invention. The basic concept of packet switching provides a highly flexible, shared network in support of interactive computer communications in a WAN. Prior to 20 packet switching, users spread over a wide area with only infrequent traffic had no cost effective way of sharing computer resources. Asynchronous communications are bursty in nature and send only small amounts of data with lots of idle time between bursts. Having dedicated lines for such communication is a waste of bandwidth and expensive. Packet switching solved those problems by providing connections as needed which were billed on the 25 number of packets transmitted. Packet switching also improved the error performance. Typically a packet switched network uses a dial up connection to a packet switching node. Once the connection to packet switching node is made, a control packet is sent to establish the session with the target host. The control packet is forwarded across the most direct link that is available in a number of hops between nodes. The target host responds with a control 30 packet sent back to the source to establish the session. Each packet is numbered sequentially and transmitted. ISDN is an entirely digital suite of dial-up data communication services delivered over the twisted pair local loop. ISDN lines have B channels that carry

PATENT

information, D-channels that carry data for signalling and control, H-channels that carry high speed data for bandwidth intensive applications. It has been a commercial failure.

Frame relay networks were first deployed in the mid 90's and is somewhat like packet switching in that each frame is individually addressed. Frame relay makes use of special switches and a shared network of very high speed. Unlike packet switching, frame relay supports the transmission of virtually any computer data stream. Frames are variable in length up to 4096 bytes. Frame relay is data oriented and does not support voice or video very well. As is the case for X.25 packet switching, frame relay overhead is high and delays in transmission are expected. Further, network congestion can result in loss of data. Although frame relay networks appear to the customer to be one-hop networks, they really are not one hop nets. There are many links between multiple Central Office (CO) switches inside the typical frame relay cloud. Each hop adds latency and the possibility of running into bandwidth congestion. Further, frame relay networks cannot cross telephone company boundaries so all sites on a frame relay WAN must be using the same frame relay provider, i.e., it is not possible for some sites to be coupled to AT&T frame relay COs and other sites to be coupled to MCI or Sprint COs. Every frame has a DLCI code in the header that identifies the customer and the virtual data path through a particular telephone company for the traffic. Therefore, it is not possible to mix frames with different DLCIs because different telco DLCIs have different formats and that will disrupt the routing process for such frames through the CO switches. If two locations on a frame relay network cannot both be served by the same frame relay provider, a second frame relay cloud must be built and the two clouds connected together by two routers at some common location that can be coupled to both clouds with the two routers coupled together by a local area network.

Cell switching has been conventionally thought to be the future of data communication networks. Cell switching encompasses both ATM networks and Switched Multimegabit Data Service (SMDS). Data is organized into cells of fixed length of 53 octets and are shipped across high speed facilities and switched in high speed, specialized switches. ATM is primarily data oriented, but it supports voice and video effectively. Cell switching is high cost and has high overhead and suffers from a lack of fully developed standards. ATM networks are also not widely commercially available yet.

The problem with all these approaches is that they are expensive with recurring telephone company charges.

PATENT

The internet as a backbone has recently loomed as a possibility for implementing wide area networks and lowering the cost. However, there are several problems with using the internet as a WAN backbone. Generally, these problems all relate to quality of service. Quality of service has to do with both errors in transmission as well as latency. Latency or 5 delay on critical packets getting from source to destination can seriously slow or disrupt operations of computer systems. Latency can also destroy the efficacy of streaming video, streaming audio and streaming multimedia product and service delivery by causing visible and/or audible gaps in the presentation of the program encoded in the data to the user or freezes. This can be very distracting and undesirable in, for example, video conferences, 10 video-on-demand, telephone calls etc. Latency is also a problem when large documents are being downloaded because it slows the process considerably. Latency arises from multiple hops between nodes on the internet coupling the source to the destination.

Prior art attempts to use the internet as a backbone did not control the number of hops and available bandwidth in the data path from source to destination. As a result the 15 number of router hops along the route and the lack of available bandwidth precluded the use of the internet as a viable private network backbone alternative. ISP's built local businesses without regard to the customers regional, national or international presence as their objective was only to offer LOCAL internet access. This resulted in attempts to use the internet as an alternative private network backbone of routes that may have few hops or 20 many hops. Routes that may have inadequate bandwidth for the worst case bandwidth requirement of a WAN were sometimes picked and that resulted in failure. This uncontrolled hop count, and lack of control of the data paths and the available bandwidth and the resulting latency caused problems in implementing WANs on the internet.

Another major problem with using the internet as a backbone is security or privacy. 25 Since the internet is a public facility, private and sensitive data transmitted over the internet is subject to snooping.

Thus, there has arisen a need for a system which can use the internet as a WAN backbone to help decrease the costs of data transport while not suffering from the aforementioned latency, privacy and bandwidth availability problems.

30

Summary of the Invention

The wide area network technology described herein (referred to as AlterWAN™ network) is an alternative wide area network that uses the internet as a backbone with any telephone company providing the local loop connection to the first participating ISX/ISP and any telephone company providing a local loop connection from the endpoint participating ISX/ISP to the destination router. This greatly reduces monthly costs to customers and removes the frame relay restriction that the same telephone company must provide all data paths including the local loops at both ends. High quality of service is maintained by mimicking the "one hop" private network structures of prior art frame relay and point-to-point networks. Any WAN that uses the internet as a backbone and mimics the "one hop" structure of private frame relay and point-to-point networks by any means is within the genus of the invention.

A key characteristic that all species within the genus of the invention will share is a tuning of the internet network routing process by proper ISX selection to reduce the hop count thereby reducing the latency problem that has plagued prior failed attempts to use the internet as a WAN backbone.

Another key characteristic that all species within the genus of the invention will share is the transmission of secure encrypted data along preplanned high bandwidth, low hop-count routing paths between pairs of customer sites that are geographically separated.

The encrypted AlterWAN data is sent through a high bandwidth, dedicated local loop connection to the first participating AlterWAN ISX/ISP facility. There, the AlterWAN packets are routed to the routers of only preselected ISX facilities on the internet. The preselected ISX/ISP facilities are ones which provide high-bandwidth, low hop-count data paths to the other ISX/ISP facilities along the private tunnel. The routers of these participating ISX/ISP facilities are specially selected to provide these high-bandwidth, low hop-count data paths either by their natural routing tables or by virtue of special routing tables that these ISX/ISP providers establish to route AlterWAN packets through high-bandwidth, low hop-count paths and route other internet traffic along other paths. For example, if a customer site in San Jose needs to have AlterWAN service to another site in Tokyo, a "private tunnel" is built in each direction through the internet and two dedicated local loops, one at each end are established to connect the two customer sites to the first and last participating ISX providers in the private tunnel. Data security is implemented by the

PATENT

use of conventional or custom firewall/VPN technology. At each customer site, a firewall/VPN device is configured to securely encrypt the payload of each AlterWAN packet to be sent through a "private tunnel" to the far end customer site where the payload is decrypted. Using conventional firewalls, the encryption method and the encryption keys used at both ends for transmissions in both directions are the same. However, the invention also contemplates using one encryption algorithm and encryption key or keys for downstream transmissions and another encryption method and different key or keys for the upstream direction. This method may require the use of custom designed firewalls.

5 Whichever method is used, the firewalls at both ends use the same encryption method and key or keys for encryption of packets at the source and decryption of them at the destination by predetermined configurations that are programmed into the firewalls. Only packets identified at the source end firewall with a destination IP address at the other end of an AlterWAN "private tunnel" have the payload of the packet encrypted before being sent. Once they are encrypted, they are sent across the preplanned route to the destination where the far end firewall recognizes the IP address of the packet as being addressed to it. Only those 10 packets are decrypted and transmitted to the device to which they are addressed and other packets that are not AlterWAN packets are either rejected or routed to some other device 15 which is not part of the AlterWAN network.

In other words, the quality of service problem that has plagued prior attempts is solved by providing non-blocking bandwidth (bandwidth that will always be available and will always be sufficient) and predefining routes for the "private tunnel" paths between points on the internet between ISX facilities. Participating ISX facilities agree to provide non-blocking bandwidth between their sites. By having private tunnels to each location of a worldwide company for example, an engineer in San Jose can connect directly to a LAN at a 20 branch office in Paris and "see" on his/her computer's desktop all the shared items on the Paris LAN such as various servers, printers etc.

25 This preplanning of the routing path causes traffic from AlterWAN™ customers to be transmitted quickly and without delay from end to end and not experience delays due to lack of bandwidth or excessive hop count. Because the packet payload is encrypted, the data is secure during its transport across the internet through the "private tunnel". The 30 AlterWAN™ network design minimize the number of hops each AlterWAN™ network packet experiences in its travel from source to destination thereby reducing latency by causing

PATENT

AlterWAN™ network traffic to be routed only over high bandwidth lines coupling participating ISX/ISP providers. Recently, there has been a large amount of building of ISX internet providers having fiber optic data paths to other providers to provide large amounts of bandwidth. Typically, one or both of the routers at the source and destination of the
5 AlterWAN™ network can be co-located at the first ISX.

The privacy problem is overcome by firewalls provided in the AlterWAN™ network at every customer premises which are encrypting firewalls (preferred firewalls are commercially available from Netscreen). Every outgoing AlterWAN™ packet (AlterWAN packets are those packets which are encrypted and are transmitted along predefined routes through the internet in "private tunnels") is encrypted by the firewall at the source using a preconfigured encryption algorithm although any encryption algorithm such as conventional DES encryption that uses a key will suffice. The encryption process requires the preprogramming of "private tunnel" identities and the associated encryption and decryption keys. The "key" is used by the firewall/VPN device for encryption and decryption of the packet payload. Keys are preassigned for each "private tunnel" and are generated by the firewalls at each end from one or two passwords that are programmed into the firewall when the private tunnel is set up. Encrypted packets are routed over predefined paths. Packets intended for the general internet are not encrypted and are passed out to the first ISX to find their way through the internet by the normal routing process. Each packet that is intended for a predefined private tunnel is encrypted and sent out through a dedicated high bandwidth local loop to the first ISX. From there it is routed along a predefined route established by proper selection of ISX providers.
10
15
20

The key can remain the same over time or change, but no packet encrypted without the current key for a particular tunnel can be decrypted at the destination. The keys are never transmitted along the tunnels. They are configured into the firewalls by prearrangement at each end. Each tunnel has a different key.
25

A "private tunnel" is defined as the data path through the internet from the source firewall to the destination firewall through the predefined, low hop count, high bandwidth path. The private tunnel is established by proper selection of ISX providers. This is done by studying the normal routing paths used by all the ISX providers between a pair of customer sites to be connected by the tunnel. Then ISX providers which normally route along high bandwidth links with a minimum hop count are selected to participate. When AlterWAN
30

PATENT

packets reach these ISX providers, the normal routing that occurs there results in the AlterWAN encrypted packets travelling along a high bandwidth low hop count path.

The ability of firewalls to encrypt and decrypt is known and commercially available and is simply being used in the AlterWAN network. Browsers at workstations at customer 5 AlterWAN sites however can be pointed to any website on the internet and can send and receive packets to and from those sites without restriction. Those packets are referred to herein as conventional packets, and they get to their destinations by conventional internet routing and do not pass through the private tunnels created by the AlterWAN data structures.

10 The AlterWAN data structures really are just IP addresses and associated data in the firewalls and routers along the tunnel that cause the packets to travel the low hop count path. The AlterWAN data structures will vary from customer to customer depending upon which sites are to be linked and the locations and IP addresses of the participating ISX/ISP providers through which the hops of the private tunnel will pass.

15 Finally, all species in the genus of the invention will solve the bandwidth bottleneck that has plagued prior attempts to use the internet as a WAN backbone. This is done by implementing AlterWAN™ routing strategies. An AlterWAN data path extends from a source router (having a channel service unit to interface between the packet world of routers to the physical and media access control and/or signalling protocols of the telephone line) through a sufficiently high bandwidth dedicated local loop line to the first participating ISX or 20 Internet Service Provider (ISP) that is a participating provider of AlterWAN™ network services. From there it extends along a data path between other participating ISX providers along a data path which is guaranteed to have sufficient bandwidth to be able to handle the worst case bandwidth consumption of the customer. In the claims, such an ISX or ISP provider is referred to as a "participating ISX/ISP". All the ISX or ISP facilities that are 25 participating in the AlterWAN™ network structure have fiber optic or other high bandwidth data paths such as OC3 or OC12 data paths available to them to send data to other ISX/ISP facilities that are participating in the AlterWAN™ network. It is these high bandwidth links which are referred to as "core bandwidth" between participating ISX/ISP facilities. It is this core bandwidth over which AlterWAN™ "private tunnel" traffic is routed on the 30 internet backbone.

The dedicated lines from the source router at the customer premises to the nearest participating ISX/ISP is typically T1 class or better in bandwidth, but it only needs to have

PATENT

two characteristics: it must be dedicated and not dialup; and, it must have sufficient bandwidth capacity to handle the worst case bandwidth consumption of the particular client facility it serves. Such a line is referred to in the claims as a "dedicated line". Thus, the dedicated lines from the source router to the nearest participating ISX/ISP may also be DSL or fractional T1.

5 The "participating ISX/ISP" to which the "dedicated line" couples may not be the nearest ISX/ISP since it is a rule of the AlterWAN™ network to only choose ISX/ISP facilities that restrict the loads in their data paths so as to have large amounts of spare bandwidth capacity. For example, AboveNet typically has loads of 50% or less in their high 10 bandwidth data paths to other ISX facilities. Therefore, AlterWAN™ network species will all have their dedicated connections to ISX/ISP facilities that have lots of spare bandwidth capacity and definitely more than the anticipated worse case bandwidth consumption of the customer so there is never a bandwidth bottleneck even if that ISX/ISP facility is not the closest facility. Although the local loop costs will be higher in such situations, the savings 15 by using the internet as a backbone without quality of service problems will greatly outweigh the burden of higher local loop costs.

20 The use of the dedicated lines to the nearest participating ISX/ISP and selection of only ISX/ISP facilities that limit the traffic in their data paths so as to have a great deal of spare capacity are the two characteristics of all AlterWAN™ network species which solve the prior art bandwidth bottleneck problems.

25 The above described structure controls the three major unpredictability factors that have frustrated prior workers in the art who have attempted to use the internet to implement WANs: hop count, bandwidth availability, and latency. The advantages of the AlterWAN™ network structure and operation are: large savings in Telco charges.; clean implementation of security not requiring PC or workstations to load special client software; use of ISX core internet bandwidth with sufficient bandwidth available for worst case 30 scenarios and with a robust fault tolerant infrastructure; the ability to offer full or partial turn-key solutions to WAN needs; local loops may be a mix of different services and speeds from different providers; an apparent one hop route to each location; customer access to local router and firewall; both public and private IP addressing can be used; communications are secure through secure tunnels using encrypted packets; and not need to rely on quality of service software processes at either end to get data, voice and video through since the

PATENT

AlterWAN network controls hop count, latency and bandwidth availability inherently by its structure and operation.

Brief Description of the Drawings

5 Figure 1 is a block diagram of a WAN using the internet as a backbone according to the genus of the invention.

Figure 2 is a block diagram of the actual hardware used in a typical AlterWAN network.

Figure 3 is a logical view of an AlterWAN private tunnel.

10 Figure 4 is a block diagram of a typical AlterWAN network for a U.S. headquarters coupled to several international sites.

Detailed Description of the Preferred and Alternative Embodiments

15 Typically 60-80% of wide area network costs over a five year period are recurring telephone company charges for their frame relay and point-to-point networking services. These charges break down into: local loop charges to have the connection; a distance charge depending upon the distances between nodes; and, a bandwidth charge for the minimum bandwidth the customer specifies. These costs can typically be drastically reduced by using the internet as a WAN backbone, but only if the latency and other quality of service problems that have plagued prior art attempts can be solved. These costs can be drastically reduced over frame relay and point-to-point networks even if extra costs of crossing telephone company boundaries are not incurred. The AlterWAN™ network of the invention does not have any telephone company boundary problems to overcome.

20 Referring to Figure 1, there is shown a block diagram of a wide area network species within the genus of the wide area networks using the internet as the backbone with controlled, small hop count, reduced latency and adequate bandwidth for the worst case scenario. A work station 10 (or server or any other peripheral) is typically coupled to an encrypting/decrypting firewall 12 by a local area network represented in this case by a LAN hub or switch 14. The work station 10 or other device may also be coupled to the firewall 12 by a dedicated line in alternative embodiments, and there may be more than one workstation or other device coupled to the firewall 12 either by LAN 14 or by individual dedicated lines. The preferred firewall is manufactured by Netscreen, but any

PATENT

encrypting/decrypting firewall that uses a customer defined key to encrypt each AlterWAN™ packet that has an IP destination address at the end of an AlterWAN private tunnel will suffice.

The function of the firewall, in one embodiment, is to receive and encrypt downstream packets addressed to nodes at the destination site on the AlterWAN network and to receive conventional internet protocol packets (hereafter IP packets) addressed to some other IP address on the internet and distinguish them from AlterWAN packets and not encrypt them. Both AlterWAN and conventional IP packets are sent to the firewall from the workstation 10 or other peripherals at customer site 1, shown in dashed lines at 20. One function of the firewall 12 (and the corresponding firewall 40 at the destination) is to distinguish between AlterWAN packets and conventional IP packets. AlterWAN packets are those packets which are addressed to destinations at the end of an AlterWAN private tunnel. Conventional packets are IP packets addressed to any other IP address other than an address at the other end of an AlterWAN private tunnel. The firewall at each end of a private tunnel are configured to encrypt AlterWAN packet payloads and send them to a router at the location of firewall from which they are converted to a suitable signal format for transmission on a dedicated local loop connection and transmitted to the first ISX/ISP provider along predefined high bandwidth, low hop-count private tunnel through the internet.

Conventional IP packets are not encrypted and are sent to the router and on the same dedicated local loop connection to the first participating ISX/ISP where they are routed to their destinations without using the private tunnel high bandwidth, low hop-count route. The firewalls make this distinction by examining the packet headers and using the destination address information and one or more lookup tables to determine which packets are AlterWAN packets addressed to nodes on the AlterWAN network and which packets are addressed to any other IP address outside the AlterWAN network.

More specifically, at each end of a private tunnel, a packet addressed to any of the IP addresses of devices at the other end of a private tunnel are recognized as packets that need to be converted to AlterWAN packets, encrypted by the firewall and encapsulated in another IP packet having as its destination address the IP address of the untrusted side of the firewall at the other end of the private tunnel. The composite AlterWAN packet is comprised of the encrypted original IP packet with an AlterWAN packet header which has as its destination address the IP address of the untrusted side of the destination firewall. At the firewall at the

PATENT

other end, these incoming AlterWAN packets will be recognized because their destination addresses match the IP address of the untrusted side of the firewall. The firewall then strips off the AlterWAN packet header of the encapsulating packet and decrypts the original IP packet that was encapsulated using the same encryption algorithm and key or keys that were used to encrypt it. The decrypted packet then has an IP packet header which has a destination address which matches the IP address of some device on the LAN on the trusted side of the destination firewall. The decrypted packet is then put on the destination LAN and makes its way to the device to which it was addressed.

The main function of the firewall is to encrypt the payloads of only AlterWAN packets with customer defined key or keys which are configured to be the same in the firewalls at both ends of the tunnel. In the preferred embodiment, commercially available firewalls are used which are configured to use the same encryption algorithm and encryption keys at both ends of each tunnel for packets travelling in either direction along the tunnel. However, in alternative embodiments, firewalls may be used which use one encryption algorithm and set of one or more encryption keys for packets travelling in one direction along the tunnel and another different encryption algorithm and/or a different set of keys in the firewalls at each end of the tunnel for packets travelling in the opposite direction. The corresponding firewall/VPN device at the tunnel far end must be programmed with the exact same key used to encrypt the packet to decrypt the packet. The encrypted packet is tested with the local key to decrypt the packet. If a match exists, the packet is decrypted and allowed through the firewall/VPN device. If not, it is discarded. Many firewalls set the encryption method and key the same for both directions of a private tunnel. In the event a firewall/VPN device implements a private tunnel by using a different encryption method and/or key for each half of a private tunnel, and that both firewall/VPN devices are configured properly, they may be implemented and used in an AlterWAN network solution. The key can be the same for all AlterWAN packets over time or it can change over time. Any encryption algorithm capable of doing this will suffice. Any conventional IP packets are not encrypted by the firewall and are simply forwarded to a router such as source router 18 or destination router 42.

The firewalls 12 and 40 are typically coupled by another local area network line to a router at the source or destination site. For example, firewall 12 is coupled by LAN line 16 to a router 18 at customer site 1, and firewall 40 is coupled by a LAN line 44 to destination router 42. Routers 18 and 42 each function to route AlterWAN and conventional IP packets

PATENT

5 differently. Both routers 18 and 42 route any AlterWAN packet into a "private tunnel" of a dedicated high bandwidth local loop data path 22 which guides these AlterWAN packets to the first participating ISX/ISP 24 in the AlterWAN™ network. The first and last participating ISX/ISP providers also have channel service units represented by boxes 23 and 25. Any conventional IP-packets are also routed into dedicated data path 22, but these conventional data packets are not part of the AlterWAN private tunnel because their destination addresses are not the address of the destination at the other end of the tunnel. Each of routers 18 and 10 42 includes a channel service unit, as shown at 19 and 21. These channel service units convert the digital data of the packets into signals suitable for transmission on whatever type of dedicated local loop signal path 22 and 46 are selected. The local loop dedicated signal paths 22 and 46 do not need to be the same type of signal path at both ends so long as suitable channel service units or cable modems are selected for CSUs 19 and 21.

15 The dedicated line 22 is typically a T1 class, partial T1 or DSL line or better with adequate bandwidth in both directions to meet the worst case bandwidth consumption scenario. DSL lines are typically not preferred since they typically only have about 640 Kbps bandwidth upstream to the CO even though they have 1.544 Mbps downstream or better. There are however some ADSL variations with up to 5 Mbps upstream and 51.84 Mbps downstream from the CO to the customer sites. One variant of ADSL supports 15 Mbps upstream and 155 Mbps downstream, but the customer ADSL modem must be within 500 meters of the central office so such a line is highly impractical unless the AlterWAN 20 customer site is virtually at the CO. Since the AlterWAN™ network is bidirectional and must have sufficient bandwidth on all data path segments thereof to meet the worst case scenario, DSL lines typically cannot be used unless the worst case scenario does not exceed the DSL line upstream bandwidth specification. Also, for DSL lines, the CO must be within 25 about 2 miles (0.6 to 1.2 miles for the higher speed variants) from the customer site and this restriction can rule out their use if a deal with a participating ISX/ISP within that range cannot be made.

30 Each of routers 18 and 42 have a channel service unit (not separately shown) built into the router (or external). The function of these channel service units in the local loop is to electrically and physically convert the (LAN) ethernet data to the signalling protocols and signal format of the telco on whatever dedicated lines 22 and 46 are chosen. The dedicated lines can be different (telephone lines or hybrid fiber coax of a CATV system or digital cable

PATENT

or satellite bidirectional links) and can be provided by different vendors. For example, if the dedicated line 22 is a T1 line the channel service units converts the AlterWAN packet data into signaling compatible with the chosen telco and transmission of that data to the matching CSU/router at the other end of the local loop where the signal is converted back to a format acceptable for processing by the router at the ISX. If the dedicated line is the hybrid fiber coaxial digital cable of a CATV system using frequency division multiplexing or code division multiplexing or discrete multitone modulation, the channel service unit modulates the ethernet onto the proper FDMA carriers or spreads the spectrums for transmission across the "local loop" with the spreading codes dedicated to the AlterWAN connection. This interfacing is bidirectional between the signal formats and protocols on dedicated lines 22 and 46

Routers 18 and 42 are the translators of the AlterWAN™ network private tunnel. The routers translate from ethernet protocol to the telco protocol on the dedicated lines 22 and 46. Other conventional IP packets that reach router 18 are routed along the same physical path and the dedicated lines but really are on a different logical path. Their payloads are not encrypted and they are not sent through the "private tunnels". AlterWAN packets addressed to different destinations will be routed into the proper private tunnels of the AlterWAN network set up for those destinations. In some embodiments, conventional IP packets will be blocked by router 18 from entering the private tunnel or any other logical channel of the dedicated lines 22 and 46. Data path 26 leaving router 18 is a DMZ path and is optional. Likewise, destination router 42 includes a DMZ port 27. The DMZ path can be any other data path that is not part of the AlterWAN network, and is typically where mail servers reside.

One of the side effects of having the high speed dedicated line 22 is that workstations at the client facility 1 (and the client facility at the other end of the WAN) can also have high speed internet access to other websites that have nothing to do with the AlterWAN solution without a separate connection. The AlterWAN traffic on dedicated line 22 shares this transport with non-AlterWAN traffic so it is important that the bandwidth on this dedicated local loop meet the aggregate needs of both AlterWAN traffic and conventional traffic. As part of this process, packets that are not AlterWAN packets are recognized by the firewall by looking at the addressing information in packet header information and are not encrypted. Conversely, packets that appear to the firewall to be addressed to nodes in the AlterWAN

PATENT

network have their packet payloads encrypted. All the packets are then sent to the source router 18 (or destination router 42) which routes them. Conventional packets get routed on dedicated line 22 other than the AlterWAN private tunnel to the first participating ISX/ISP 24. At the first ISX/ISP 24 in the AlterWAN network, these conventional packets get routed out one of the data paths represented by lines 27 through 36 that couple router 24 to the rest of the internet. This provides high speed access to other web pages and websites and e-mail services as a byproduct of the AlterWAN hardware and software processing.

AlterWAN packets get routed at the first ISX/ISP 24 into a high bandwidth data path 50 to the next participating ISX/ISP 48 in the AlterWAN network. Data path 50 is selected for the AlterWAN packets by the preselected ISX/ISP and peer level predefined routing between participating ISX/ISP's. This allows AlterWAN traffic to be transported between locations utilizing the naturally existing routes but those routes are selected so as to be high bandwidth and low hop count. Each router in the participating ISX/ISP facilities connects and communicates in the same fashion. AlterWAN networks, by design, require selection of the ISX/ISP partners for any given network based on many factors including the ease of implementation by utilizing naturally occurring or other existing high bandwidth, low hop count routes. AlterWAN designers pretest these routes by performing a minimum of a ping test and traceroute test to verify the path data that AlterWAN packets will take through the private tunnel that is to be implemented as an AlterWAN connection. AlterWAN partners do not normally need to add special routes, but implementing AlterWAN network designs that follow existing known paths does not preclude the addition of special routing from time to time as needed to afford better routing. By such a process, an AlterWAN network does not require each participating ISX/ISP to make alterations to their equipment for each "private tunnel" created but rather transparently utilizes the high bandwidth peer level connections between ISX/ISP's. However, the invention does not preclude use of ISX/ISP providers who have altered their routing tables so as to insure that AlterWAN packets get routed along high bandwidth, low hop-count data paths while non-AlterWAN packets get routed along other data paths. Participating ISX/ISP's are selected in part based on their ability to use these natural routes to form low hop count connections between the ends of an AlterWAN private tunnel or by entering into a special deal with one or more other participating ISX/ISP's to implement special peering arrangements and/or routing between each other to allow only

PATENT

AlterWAN traffic to use these special low hop count high bandwidth connections forcing non AlterWAN traffic to follow other natural routing that does not provide the bandwidth and or hop counts that meet the AlterWAN requirement.

In the example of Figure 1, only three participating ISX/ISP providers are shown at 5 24, 48 and 54. The high bandwidth paths are the naturally occurring data paths that result from the routing tables in the participating ISX provider routers. These data paths are represented by lines 50 and 56. The private tunnel between customer site #1 at 20 and customer site #2 at 58 is implemented by the dedicated lines 22 and 46 and the high bandwidth data paths 50 and 56 selected for AlterWAN packets by the routing tables in 10 participating ISX/ISP providers 24, 48 and 54.

When AlterWAN packets from customer site #1 reach endpoint ISX/ISP router 54, they are routed onto dedicated line 46 to the channel service unit of destination router 42. The destination router 42 recovers and reassembles the ethernet packets and outputs them to firewall 40. Firewall 40 decrypts all AlterWAN packets with its local matching key preconfigured on the firewall/VPN device and formats them to the LAN protocol. It then forwards them to the destination LAN hub or switch 60 where they are sent out on LAN 62 addressed to whatever peripheral 64, 66 or 68 to which they are destined. AlterWAN packets from any of these peripherals addressed to any of the peripherals at customer site #1, 20, are encrypted by firewall 40 and are routed back through the private tunnel to site 20 where they are decrypted by firewall 12 and forwarded to LAN hub or switch 14 and sent out on LAN 70 to whatever peripheral at site 20 to which they are addressed.

FIREWALL AND TUNNEL SETUP

The firewalls 12 and 40 can be any commercially available firewall with the ability 25 to create a virtual private network. The firewalls serve two general purposes: they provide general security from unwanted access to the AlterWAN customer LAN network; and they provide private encrypted tunnels between a known set of sites even though the internet is a public facility. Each customer's AlterWAN network will be different based upon their needs in terms of the type and bandwidth of dedicated lines used and the private tunnel data paths 30 set up through the participating ISX/ISP providers between customer sites.

The interfaces of a firewall consist of an untrusted WAN interface, one or more trusted IP interfaces to dedicated lines or LAN drop lines, and a DMZ interface (if available).

PATENT

These three interfaces are illustrated at 72, 74 and 76, respectively, in Figure 2 which is a block diagram of the actual hardware configuration of a typical AlterWAN network. The untrusted or WAN interface is used to interface to the ISX/ISP premises router of the public internet, optionally through a customer premises router 18 or 42. The IP trusted interface 5 interfaces to the customer's private local area network 70 or 62 (or to dedicated lines to each peripheral in some embodiments). The DMZ interface (optionally available on some firewalls) is used to configure a separate network where devices that may need both public and private access typically are placed including WEB servers and e-mail servers.

Every LAN and WAN interface at both the customer premises and the ISX/ISP in 10 Figure 2 needs to be configured with IP addresses. The exception to this would be any LAN using a protocol different than ethernet IP such as Token Ring. In such case the proper networking and conversion equipment would be required. Each interface to be configured in general includes: an IP address, for example 204.123.111.150; a network mask, for example 255.255.255.0; and a default gateway, for example 204.123.111.1. The addressing for each interface is either supplied by the ISX/ISP or by the customer. The telephone (or cable system operator) high bandwidth dedicated lines 22 and 46 need to be in 15 place and operational in addition to the configurations mentioned above to complete the AlterWAN structure.

Tunnels and encryption methods vary between manufacturers of firewalls and 20 virtual private network (hereafter VPN) equipment. This limits the ability to mix products from different manufacturers within a specific customer's AlterWAN setup because the firewalls/VPN process at each end of each tunnel must use the same encryption algorithms so AlterWAN packets can be properly encrypted and decrypted. If however, all firewalls from all manufacturers can be modified to use the same encryption algorithm, then 25 firewalls/VPN processes from different manufacturers can be mixed and matched. The VPN processing hardware and software to encrypt and decrypt AlterWAN packets can be integrated into the firewall or external to it.

A virtual private network tunnel requires the following basic components and data 30 structures at each end of the tunnel. There must be a virtual private network process running on a VPN processor (can be the same processor as the firewall processor) or external to a firewall on each end of the private tunnel. The untrusted address of the far end VPN untrusted WAN interface must be configured in the VPN configuration data structure at

PATENT

each end including a mnemonic label, an IP address and a network mask. The VPN configuration data structure at each end must also include a mnemonic label, an encryption key, an encryption type, an encryption password, and the gateway IP address of the far end firewall untrusted or WAN interface. Only when a VPN pair configured in this manner 5 exists with one VPN on each end of a proposed tunnel, and the participating ISX/ISP providers route a path between the two endpoints over high bandwidth links with a minimum number of hop for AlterWAN packets, does the private tunnel actually exist. Once the tunnel is created, all the conventional internet routers and uncontrolled number of hops and 10 uncontrolled latency that they create for non AlterWAN packets virtually disappear for AlterWAN packets. The AlterWAN data path logically appears to be a direct point-to-point connection between the two sites at opposite ends of the tunnel as shown in Figure 3.

Private tunnels are defined for each customer based upon the needs of that customer. This is done by identifying a set of known participating ISX/ISP locations through which the 15 number of known hops caused by their routers is minimized. All locations on the internet outside this known set of sites and the associated networks are assumed to be general internet sites to which conventional IP packets can be directed.

The only real difference between a conventional IP packet and an AlterWAN packet is that the payload of the AlterWAN packet is encrypted. Conventional packets have no 20 encryption performed on the packet payload and are routed to the default gateway IP address of the participating ISP/ISX.

The firewalls at each end of each private tunnel prevent any unauthorized user from 25 accessing the private LANs of AlterWAN customers. The tunnels in each firewall have configuration data that only allows specific user traffic access to the private tunnels. Traceroutes to any address outside the tunnel show all router hops for conventional packets while traceroutes to any address inside a private tunnel show only private tunnel hops for 30 AlterWAN traffic. The establishment of a private tunnel enables users at a first customer site to appear to be directly connected to a LAN at another site of the customer so that all the shared resources on the other end of the tunnel appear on the desktops of the workstations at the first site. Most of the participating providers in AlterWAN structures are ISX providers. This eliminates the numerous hops customers typically incur in dealing with local ISPs for wide area networking. By picking participating ISX providers that have high bandwidth lines that are not fully utilized, the bandwidth availability problem of using the

PATENT

internet as a WAN backbone is solved. Numerous ISX providers now offer 1-hop connections to major cities in the U.S. and throughout the world. The AlterWAN network structure takes advantage of this fact by selecting the ISX/ISP's that form the shortest path(s) between the set of customer sites that need to communicate. Through this design and selection process, 5 the natural routes that stitch together these high bandwidth single hop lines with dedicated high bandwidth local loops to geographically separated customer sites to create a private tunnel through the internet between any two customer sites to provide frame relay quality service at substantially less cost.

Frame relay prior art WANs were considered highly desirable because they establish 10 permanent virtual circuits with known paths having known bandwidth. The internet has not been able to provide a similar solution in the prior art. The AlterWAN network structure changes that by creating virtual private circuits or tunnels through the internet using only lines that are known to have sufficient bandwidth to carry the worst case load and by 15 minimizing the number of hops by using primarily ISX providers. Prior attempts to use the internet for WANs have failed because the data paths were not controlled, the bandwidth was oversubscribed or in any fashion insufficient causing unacceptable latency and delays. This caused unpredictable latency which is very undesirable for multimedia video and audio traffic. Only light users with small amounts of non time sensitive data were able to use the internet successfully as a WAN. The AlterWAN network structure uses a set of known high 20 bandwidth, usually fiber optic, links between major domestic and international cities and couples these data paths with dedicated point-to-point or frame relay circuits run locally from the "nearest" participating ISX/ISP (sometimes it is not the physically nearest ISX but is the nearest ISX with a high bandwidth line to a key city that is participating) to the customer site. The unique aspects are forcing the participating routers to stitch together 25 known high bandwidth data paths with a minimum number of hops to high bandwidth dedicated local loop connections and encrypting all AlterWAN traffic for privacy.

Figure 4 is block diagram of a typical AlterWAN network for an international corporation with multiple international locations in the U.K., Germany, France and Japan with a headquarters in the U.S. Suppose one of workstations 78 through 82 on LAN 84 in the 30 U.K. site 96 wants to have access to server 86 on LAN 88 at the U.S. headquarters. Workstations 78 generates an IP packet that gets encapsulated into an Ethernet or other LAN packet addressed to the firewall 90. The firewall looks up the IP address in its tables and

PATENT

determines that the packet is addressed to an AlterWAN IP address in the U.S. headquarters. It then encrypts the payload portion of the packet using the prearranged key for the tunnel to the destination. The encrypted packet payload is sent through the "private tunnel" from the U.K. firewall 90 to the U.S. site firewall 92. Network address translation unit converts any 5 IP addresses that conflict with private IP addresses owned by some other company to one IP address on the untrusted interface given by the participating ISX. Firewalls can handle both NAT addressing and transparent addressing, but that is not relevant to the invention.

After encryption, the AlterWAN packet is forwarded to router 98 at the U.K. site 96. This router examines every packet and based on the routing tables forwards packets to the 10 next ISX. In this case, the router will only receive packets from the firewall if they were not for the local LAN. At this time, AlterWAN packets and conventional IP packets are equal, but AlterWAN traffic has "designed in" efficient routing paths to the destination points with the ISX/ISP connected by dedicated local loop line 100 that couples the router to the first participating ISX provider within internet cloud 102 via a known internal or external channel service unit. The router in the first participating ISX within the internet cloud receives the AlterWAN packets and routes them along the predetermined private tunnel data 15 path that has been preplanned to use the natural routing table (unless a special case requires additional special routes). This process continues at each router of each ISX along the private tunnel to the U.S. site 106. The last participating ISX along the private tunnel is represented by switch 104. This switch has all AlterWAN packets destined for this location 20 passing therethrough and may be used to keep track of traffic levels for purposes of billing. Billing can be based on fixed monthly connections and/or billing with a base fee and usage fee. Collection of the information to generate billing on base fee plus usage is from each location requiring such.

25 From switch 104, AlterWAN packets are routed to firewall 92 at the customer U.S. site where they are decrypted and sent to router 108 which outputs the packets onto LAN 88 where they are received and processed by server 86. Non AlterWAN packets routed by switch 104 to firewall 110 are either replies to general internet activity initiated on LAN 88 or outside traffic requests intended for the web servers 114 and 116 on the firewall 30 DMZ. Any other traffic would be rejected by the firewall. These packets are not encrypted, and after the firewall 110 processes them, they are routed to a LAN hub 112 and sent from there to a web server 114 and another web server 116.

PATENT

AlterWAN packets that originate at the U.K. or one of the other international sites and are addressed to another international site never go to switch 104. Instead an IP packet originating at, for example, the U.S. site and addressed to a device on the LAN at the French site, get routed through a private tunnel that extends from the U.K. firewall 90 to the
5 French firewall 91. Thus, these packets never pass through switch 104.

Although the invention has been disclosed in terms of the preferred and alternative embodiments disclosed herein, those skilled in the art will appreciate possible alternative embodiments and other modifications to the teachings disclosed herein which do not depart from the spirit and scope of the invention. All such alternative embodiments and other
10 modifications are intended to be included within the scope of the claims appended hereto.

Appendix A is a typical list of configuration commands for the firewall at the headquarters site of a typical AlterWAN to establish a private tunnel through the internet from the headquarters to a destination site firewall including establishment of the IP address of the first ISX in the tunnel. Appendix B is a typical list of configuration commands for the destination site firewall at the other end of the private tunnel. Appendix C is a typical list of configuration commands to configure the router at the headquarters site. Appendix D is a typical list of configuration commands to configure the router at the destination site.
15

Although the invention has been disclosed in terms of the preferred and alternative embodiments disclosed herein, those skilled in the art will appreciate possible alternative embodiments and other modifications to the teachings disclosed herein which do not depart from the spirit and scope of the invention. All such alternative embodiments and other
20 modifications are intended to be included within the scope of the claims appended hereto.